DECISIONFLOW

Business Viability and Unit Economics Report

Pricing, renewal, revenue quality, cost structure, and scale economics

QUALIFY PHASE | Version 1.0 | 14 July 2026

EVIDENCE & PUBLISHING INTEGRITY

This public portfolio edition distinguishes actual evidence, assumed continuity steps, simulated case-study outcomes, targets, and recommendations. Simulated customers, commercial results, benchmarks, incidents, quotations, and later-lifecycle outcomes are illustrative and must not be interpreted as results from a live commercial product. Evidence labels are retained throughout the package to support transparent and responsible interpretation.

Document Control

Control	Definition
Document code	QUAL-03
Version / status	2.0 — Final Public Portfolio Edition — Simulated Qualification Scenario
Baseline / author	14 July 2026 — Min Thu Kyaw
Role / purpose	AI Product Manager — Business Viability and Unit Economics Report
Lifecycle continuity	Actual MVP → DEV-06 illustrative productionization → simulated v1.0 launch → qualification decision.
Evidence rule	Actual artifacts, assumed capabilities, simulated outcomes, and recommendations are labelled separately.

Actual Project Artifacts

Artifact	Reference
Low-Fidelity Prototype	View low-fidelity prototype
High-Fidelity Prototype	View high-fidelity prototype
MVP Frontend Repository	Open frontend repository
MVP Backend Repository	Open backend repository

Contents

Viability Objective

The business qualification asks whether DecisionFlow can convert product value into sustainable revenue without depending on unusually high founder involvement, hidden AI costs, or capabilities that are not part of v1.0. The analysis treats the controlled launch conversion rate as directional evidence only; it does not assume that a broader self-service funnel will perform similarly.

Commercial Cohort

Measure	Simulated result	Evidence quality
Organizations invited	20	Controlled cohort; low generalizability.
Organizations onboarded	15	Qualified, champion-led accounts.
Organizations paid at Week 12	8	Strong directional signal; small sample.
Original paid organizations retained at Week 24	7	Early renewal evidence.
Organizations expanded	2	Positive depth signal.
Organizations churned	1	Churn linked to low decision frequency.
Average paid users at launch	1.6 per organization	Supports individual packaging, not team-product proof.

Revenue and Retention

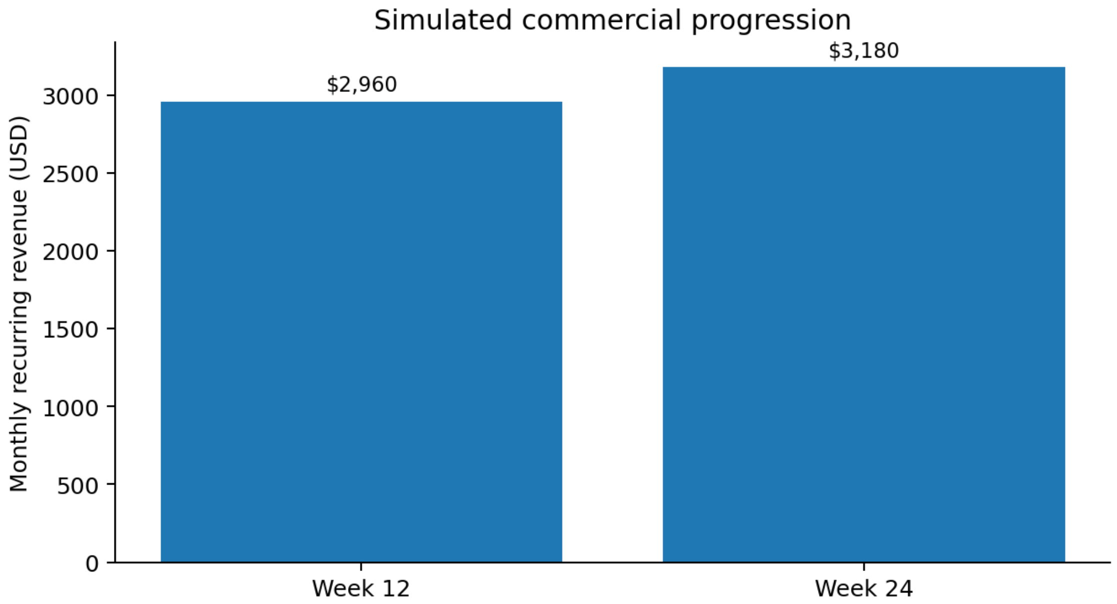

Commercial progression

Metric	Formula	Simulated result	Interpretation
Gross revenue retention	(Starting MRR − churned MRR) / starting MRR	92.6%	Positive early retention signal.
Net revenue retention	(Starting MRR − churn + expansion) / starting MRR	107.4%	Expansion offsets churn.
Week-24 ARPA	$3,180 / 7 retained paid organizations	$454	Higher than launch due to expansion.
Annualized run rate	$3,180 × 12	$38,160	Illustrative run rate, not actual revenue.
Paid cohort logo retention	7 / 8	87.5%	Promising; confidence interval is wide.

Unit Economics

Measure	Simulated / modeled value	Gate	Status
Gross margin	82%	≥80%	Pass
Variable cost per finalized decision	$0.38	≤$0.60	Pass
Support-adjusted contribution margin	47%	≥55%	Fail
Modeled blended CAC	$2,100	Directional	Unproven at scale
Gross-profit payback	6.9 months	≤12 months	Pass model
Modeled LTV / CAC	3.2×	≥3.0×	Pass model, assumption-sensitive
Median procurement cycle	24 days	≤30 days	Pass

Unit-economics warning — The strongest commercial numbers come from a controlled, guided launch. Support-adjusted contribution margin is below the scale threshold, and customer acquisition cost is modeled rather than observed through repeatable paid channels. This is not evidence for unrestricted marketing spend.

Pricing Evidence

Hypothesis	Simulated evidence	Decision
Starter fits low-volume individual users	Two paid accounts preferred lower limits and simpler pricing.	Retain a low-friction Starter package.
Professional fits frequent decision owners	Five of eight paid accounts described value and price as acceptable.	Keep Professional as primary v1.0 package.
Team package is ready	Average paid users are 1.6 and collaboration is not in v1.0.	Do not launch Team package.
Enterprise package is ready	Security reviews requested controls not implemented in v1.0.	Use design-partner discovery only.
Usage-based pricing improves fit	Heavy regeneration can increase cost but customers prefer predictable pricing.	Test capped decision allowances, not pure consumption billing.

Scenario Sensitivity

Scenario	Key assumptions	Modeled implication	Decision
Base case	82% gross margin; 4.5% monthly churn; $2,100 CAC	LTV/CAC ≈3.2×; 6.9-month payback	Conditional scale.
Retention downside	7% monthly churn	LTV/CAC falls below 2.2×	Freeze paid acquisition.
Support downside	Contribution margin stays below 45%	Growth increases service cost faster than revenue	Automate onboarding before scale.
AI cost upside	Cost per decision falls to $0.28	Gross margin improves modestly	Useful but not the primary constraint.
Expansion upside	25% of retained accounts expand	NRR increases, but collaboration proof remains absent	Expand individual Professional use, not Team claims.

Commercial Risks

- Controlled-cohort conversion is likely higher than open-market conversion because participants were selected and supported.
- Founder-led acquisition and guided onboarding hide the true cost of repeatable growth.

- Low decision frequency can create churn even when individual workflows are highly valued.
- Integration and security requests may lengthen sales cycles without improving v1.0 individual usage.
- AI provider cost and regeneration behavior can compress margins if limits and routing are not monitored.

Business Qualification Decision

Business decision — PROMISING BUT NOT FULLY QUALIFIED. Continue monetizing Starter and Professional within a controlled segment. Do not scale paid acquisition until support-adjusted contribution margin exceeds 55%, at least 15 paying organizations produce 90-day renewal evidence, and CAC is observed through a repeatable channel.

90-day business gate	Threshold
Paying organizations	≥15 with defined ICP and source channel
90-day gross logo retention	≥80%
Gross revenue retention	≥90%
Support-adjusted contribution margin	≥55%
Observed CAC payback	≤12 months
Professional-package price acceptance	≥60% of qualified buyers

Evidence discipline: replace all simulated values with observed, reproducible evidence before using this material as a real product-performance claim.